EXHIBIT 5

OpenAI   Menu

Terms of use

Updated

March 14, 2023

Thank you for using OpenAI!

These Terms of Use apply when you use the services of OpenAI, L.L.C. or our affiliates, including our application programming interface, software, tools, developer services, data, documentation, and websites ("Services"). The Terms include our Service Terms, Sharing & Publication Policy, Usage Policies, and other documentation, guidelines, or policies we may provide in writing. By using our Services, you agree to these Terms. Our Privacy Policy explains how we collect and use personal information.

1. Registration and Access

You must be at least 13 years old to use the Services. If you are under 18 you must have your parent or legal guardian's permission to use the Services. If you use the Services on behalf of another person or entity, you must have the authority to accept the Terms on their behalf. You must provide accurate and complete information to register for an account. You may not make your access credentials or account available to others outside your organization, and you are responsible for all activities that occur using your credentials.

2. Usage Requirements

OpenAI

(b) **Feedback**. We appreciate feedback, comments, ideas, proposals and suggestions for improvements. If you provide any of these things, we may use it without restriction or compensation to you.

(c) **Restrictions**. You may not (i) use the Services in a way that infringes, misappropriates or violates any person's rights; (ii) reverse assemble, reverse compile, decompile, translate or otherwise attempt to discover the source code or underlying components of models, algorithms, and systems of the Services (except to the extent such restrictions are contrary to applicable law); (iii) use output from the Services to develop models that compete with OpenAI; (iv) except as permitted through the API, use any automated or programmatic method to extract data or output from the Services, including scraping, web harvesting, or web data extraction; (v) represent that output from the Services was human-generated when it is not or otherwise violate our Usage Policies; (vi) buy, sell, or transfer API keys without our prior consent; or (vii), send us any personal information of children under 13 or the applicable age of digital consent. You will comply with any rate limits and other requirements in our documentation. You may use Services only in geographies currently supported by OpenAI.

(d) **Third Party Services**. Any third party software, services, or other products you use in connection with the Services are subject to their own terms, and we are not responsible for third party products.

3. Content

(a) **Your Content**. You may provide input to the Services ("Input"), and receive output generated and returned by the Services based on the Input ("Output"). Input and Output are collectively "Content." As between the parties and to the extent permitted by applicable law, you own all Input. Subject to your compliance with these Terms, OpenAI hereby assigns to you all its right, title and interest in and to Output. This means you can use Content for any purpose, including commercial purposes such as sale or publication, if you comply with these Terms. OpenAI may use Content to provide and maintain the Services, comply with applicable law, and enforce our policies. You are responsible for Content, including for ensuring that it does not violate any applicable law or these Terms.

(b) **Similarity of Content**. Due to the nature of machine learning, Output may not be unique across users and the Services may generate the same or similar output for OpenAI or a third party. For example, you may provide input to a model such as "What color is the sky?" and receive output such as "The sky is blue." Other users may also ask similar questions and receive the same response. Responses that are requested by and generated for other users are not considered your Content.

(c) **Use of Content to Improve Services**. We do not use Content that you provide to or receive from our API ("API Content") to develop or improve our Services. We may use Content from Services other than our API ("Non-API Content") to help develop and improve our Services. You can read more here about how Non-API Content may be used to improve model performance. If you do not want your

⬡ OpenAI

(d) **Accuracy**. Artificial intelligence and machine learning are rapidly evolving fields of study. We are constantly working to improve our Services to make them more accurate, reliable, safe and beneficial. Given the probabilistic nature of machine learning, use of our Services may in some situations result in incorrect Output that does not accurately reflect real people, places, or facts. You should evaluate the accuracy of any Output as appropriate for your use case, including by using human review of the Output.

4. Fees and Payments

(a) **Fees and Billing**. You will pay all fees charged to your account ("Fees") according to the prices and terms on the applicable pricing page, or as otherwise agreed between us in writing. We have the right to correct pricing errors or mistakes even if we have already issued an invoice or received payment. You will provide complete and accurate billing information including a valid and authorized payment method. We will charge your payment method on an agreed-upon periodic basis, but may reasonably change the date on which the charge is posted. You authorize OpenAI and its affiliates, and our third-party payment processor(s), to charge your payment method for the Fees. If your payment cannot be completed, we will provide you written notice and may suspend access to the Services until payment is received. Fees are payable in U.S. dollars and are due upon invoice issuance. Payments are nonrefundable except as provided in this Agreement.

(b) **Taxes**. Unless otherwise stated, Fees do not include federal, state, local, and foreign taxes, duties, and other similar assessments ("Taxes"). You are responsible for all Taxes associated with your purchase, excluding Taxes based on our net income, and we may invoice you for such Taxes. You agree to timely pay such Taxes and provide us with documentation showing the payment, or additional evidence that we may reasonably require. OpenAI uses the name and address in your account registration as the place of supply for tax purposes, so you must keep this information accurate and up-to-date.

(c) **Price Changes**. We may change our prices by posting notice to your account and/or to our website. Price increases will be effective 14 days after they are posted, except for increases made for legal reasons or increases made to Beta Services (as defined in our Service Terms), which will be effective immediately. Any price changes will apply to the Fees charged to your account immediately after the effective date of the changes.

(d) **Disputes and Late Payments**. If you want to dispute any Fees or Taxes, please contact ar@openai.com within thirty (30) days of the date of the disputed invoice. Undisputed amounts past due may be subject to a finance charge of 1.5% of the unpaid balance per month. If any amount of your Fees are past due, we may suspend your access to the Services after we provide you written notice of late payment.

(e) **Free Tier**. You may not create more than one account to benefit from credits provided in the free tier of the Services. If we believe you are not using the free tier in good faith, we may charge you

OpenAI

5. Confidentiality, Security and Data Protection

(a) **Confidentiality**. You may be given access to Confidential Information of OpenAI, its affiliates and other third parties. You may use Confidential Information only as needed to use the Services as permitted under these Terms. You may not disclose Confidential Information to any third party, and you will protect Confidential Information in the same manner that you protect your own confidential information of a similar nature, using at least reasonable care. Confidential Information means nonpublic information that OpenAI or its affiliates or third parties designate as confidential or should reasonably be considered confidential under the circumstances, including software, specifications, and other nonpublic business information. Confidential Information does not include information that: (i) is or becomes generally available to the public through no fault of yours; (ii) you already possess without any confidentiality obligations when you received it under these Terms; (iii) is rightfully disclosed to you by a third party without any confidentiality obligations; or (iv) you independently developed without using Confidential Information. You may disclose Confidential Information when required by law or the valid order of a court or other governmental authority if you give reasonable prior written notice to OpenAI and use reasonable efforts to limit the scope of disclosure, including assisting us with challenging the disclosure requirement, in each case where possible.

(b) **Security**. You must implement reasonable and appropriate measures designed to help secure your access to and use of the Services. If you discover any vulnerabilities or breaches related to your use of the Services, you must promptly contact OpenAI and provide details of the vulnerability or breach.

(c) **Processing of Personal Data**. If you use the Services to process personal data, you must provide legally adequate privacy notices and obtain necessary consents for the processing of such data, and you represent to us that you are processing such data in accordance with applicable law. If you will be using the OpenAI API for the processing of "personal data" as defined in the GDPR or "Personal Information" as defined in CCPA, please fill out this form to request to execute our Data Processing Addendum.

6. Term and Termination

(a) **Termination; Suspension**. These Terms take effect when you first use the Services and remain in effect until terminated. You may terminate these Terms at any time for any reason by discontinuing the use of the Services and Content. We may terminate these Terms for any reason by providing you at least 30 days' advance notice. We may terminate these Terms immediately upon notice to you if you materially breach Sections 2 (Usage Requirements), 5 (Confidentiality, Security and Data Protection), 8 (Dispute Resolution) or 9 (General Terms), if there are changes in relationships with third party technology providers outside of our control, or to comply with law

OpenAI

(b) **Effect on Termination**. Upon termination, you will stop using the Services and you will promptly return or, if instructed by us, destroy any Confidential Information. The sections of these Terms which by their nature should survive termination or expiration should survive, including but not limited to Sections 3 and 5-9.

7. Indemnification; Disclaimer of Warranties; Limitations on Liability

(a) **Indemnity**. You will defend, indemnify, and hold harmless us, our affiliates, and our personnel, from and against any claims, losses, and expenses (including attorneys' fees) arising from or relating to your use of the Services, including your Content, products or services you develop or offer in connection with the Services, and your breach of these Terms or violation of applicable law.

(b) **Disclaimer**. THE SERVICES ARE PROVIDED "AS IS." EXCEPT TO THE EXTENT PROHIBITED BY LAW, WE AND OUR AFFILIATES AND LICENSORS MAKE NO WARRANTIES (EXPRESS, IMPLIED, STATUTORY OR OTHERWISE) WITH RESPECT TO THE SERVICES, AND DISCLAIM ALL WARRANTIES INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, SATISFACTORY QUALITY, NON-INFRINGEMENT, AND QUIET ENJOYMENT, AND ANY WARRANTIES ARISING OUT OF ANY COURSE OF DEALING OR TRADE USAGE. WE DO NOT WARRANT THAT THE SERVICES WILL BE UNINTERRUPTED, ACCURATE OR ERROR FREE, OR THAT ANY CONTENT WILL BE SECURE OR NOT LOST OR ALTERED.

(c) **Limitations of Liability**. NEITHER WE NOR ANY OF OUR AFFILIATES OR LICENSORS WILL BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES, INCLUDING DAMAGES FOR LOSS OF PROFITS, GOODWILL, USE, OR DATA OR OTHER LOSSES, EVEN IF WE HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. OUR AGGREGATE LIABILITY UNDER THESE TERMS SHALL NOT EXCEED THE GREATER OF THE AMOUNT YOU PAID FOR THE SERVICE THAT GAVE RISE TO THE CLAIM DURING THE 12 MONTHS BEFORE THE LIABILITY AROSE OR ONE HUNDRED DOLLARS ($100). THE LIMITATIONS IN THIS SECTION APPLY ONLY TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW.

8. Dispute Resolution

YOU AGREE TO THE FOLLOWING MANDATORY ARBITRATION AND CLASS ACTION WAIVER PROVISIONS:

(a) **MANDATORY ARBITRATION**. You and OpenAI agree to resolve any past or present claims relating to these Terms or our Services through final and binding arbitration, except that you have

⊛ OpenAI

(b) **Informal Dispute Resolution.** We would like to understand and try to address your concerns prior to formal legal action. Before filing a claim against OpenAI, you agree to try to resolve the dispute informally by sending us notice at dispute-resolution@openai.com of your name, a description of the dispute, and the relief you seek. If we are unable to resolve a dispute within 60 days, you may bring a formal proceeding. Any statute of limitations will be tolled during the 60-day resolution process. If you reside in the EU, the European Commission provides for an online dispute resolution platform, which you can access at https://ec.europa.eu/consumers/odr.

(c) **Arbitration Forum.** Either party may commence binding arbitration through ADR Services, an alternative dispute resolution provider. The parties will pay equal shares of the arbitration fees. If the arbitrator finds that you cannot afford to pay the arbitration fees and cannot obtain a waiver, OpenAI will pay them for you. OpenAI will not seek its attorneys' fees and costs in arbitration unless the arbitrator determines that your claim is frivolous.

(d) **Arbitration Procedures**. The arbitration will be conducted by telephone, based on written submissions, video conference, or in person in San Francisco, California or at another mutually agreed location. The arbitration will be conducted by a sole arbitrator by ADR Services under its then-prevailing rules. All issues are for the arbitrator to decide, except a California court has the authority to determine (i) the scope, enforceability, and arbitrability of this Section 8, including the mass filing procedures below, and (ii) whether you have complied with the pre-arbitration requirements in this section. The amount of any settlement offer will not be disclosed to the arbitrator by either party until after the arbitrator determines the final award, if any.

(e). **Exceptions**. This arbitration section does not require arbitration of the following claims: (i) individual claims brought in small claims court; and (ii) injunctive or other equitable relief to stop unauthorized use or abuse of the Services or intellectual property infringement.

(f) **NO CLASS ACTIONS**. Disputes must be brought on an individual basis only, and may not be brought as a plaintiff or class member in any purported class, consolidated, or representative proceeding. Class arbitrations, class actions, private attorney general actions, and consolidation with other arbitrations are not allowed. If for any reason a dispute proceeds in court rather than through arbitration, each party knowingly and irrevocably waives any right to trial by jury in any action, proceeding, or counterclaim. This does not prevent either party from participating in a class-wide settlement of claims.

(g) **Mass Filings**. If, at any time, 30 or more similar demands for arbitration are asserted against OpenAI or related parties by the same or coordinated counsel or entities ("Mass Filing"), ADR Services will randomly assign sequential numbers to each of the Mass Filings. Claims numbered 1-10 will be the "Initial Test Cases" and will proceed to arbitration first. The arbitrators will render a final award for the Initial Test Cases within 120 days of the initial pre-hearing conference, unless the claims are resolved in advance or the parties agree to extend the deadline. The parties will then have 90 days (the "Mediation Period") to resolve the remaining cases in mediation based on the awards from the Initial Test Cases. If the parties are unable to resolve the outstanding claims during this time, the parties may choose to opt out of the arbitration process and proceed in court by providing written notice to the other party within 60 days after the Mediation Period. Otherwise, the remaining

OpenAI

(h) **Severability.** If any part of this Section 8 is found to be illegal or unenforceable, the remainder will remain in effect, except that if a finding of partial illegality or unenforceability would allow Mass Filing or class or representative arbitration, this Section 8 will be unenforceable in its entirety. Nothing in this section will be deemed to waive or otherwise limit the right to seek public injunctive relief or any other non-waivable right, pending a ruling on the substance of such claim from the arbitrator.

9. General Terms

(a) **Relationship of the Parties**. These Terms do not create a partnership, joint venture or agency relationship between you and OpenAI or any of OpenAI's affiliates. OpenAI and you are independent contractors and neither party will have the power to bind the other or to incur obligations on the other's behalf without the other party's prior written consent.

(b) **Use of Brands**. You may not use OpenAI's or any of its affiliates' names, logos, or trademarks, without our prior written consent.

(c) **U.S. Federal Agency Entities**. The Services were developed solely at private expense and are commercial computer software and related documentation within the meaning of the applicable U.S. Federal Acquisition Regulation and agency supplements thereto.

(d) **Copyright Complaints**. If you believe that your intellectual property rights have been infringed, please send notice to the address below or fill out this form. We may delete or disable content alleged to be infringing and may terminate accounts of repeat infringers.

>OpenAI, L.L.C.
>3180 18th St
>San Francisco, CA 94110
>Attn: General Counsel / Copyright Agent

Written claims concerning copyright infringement must include the following information:

- A physical or electronic signature of the person authorized to act on behalf of the owner of the copyright interest;
- A description of the copyrighted work that you claim has been infringed upon;
- A description of where the material that you claim is infringing is located on the site;
- Your address, telephone number, and e-mail address;
- A statement by you that you have a good-faith belief that the disputed use is not authorized by the copyright owner, its agent, or the law; and

OpenAI

(e) **Assignment and Delegation**. You may not assign or delegate any rights or obligations under these Terms, including in connection with a change of control. Any purported assignment and delegation shall be null and void. We may assign these Terms in connection with a merger, acquisition or sale of all or substantially all of our assets, or to any affiliate or as part of a corporate reorganization.

(f) **Modifications**. We may amend these Terms from time to time by posting a revised version on the website, or if an update materially adversely affects your rights or obligations under these Terms we will provide notice to you either by emailing the email associated with your account or providing an in-product notification. Those changes will become effective no sooner than 30 days after we notify you. All other changes will be effective immediately. Your continued use of the Services after any change means you agree to such change.

(g) **Notices**. All notices will be in writing. We may notify you using the registration information you provided or the email address associated with your use of the Services. Service will be deemed given on the date of receipt if delivered by email or on the date sent via courier if delivered by post. OpenAI accepts service of process at this address: OpenAI, L.L.C., 3180 18th Street, San Francisco, CA 94110, Attn: contract-notices@openai.com.

(h) **Waiver and Severability**. If you do not comply with these Terms, and OpenAI does not take action right away, this does not mean OpenAI is giving up any of our rights. Except as provided in Section 8, if any part of these Terms is determined to be invalid or unenforceable by a court of competent jurisdiction, that term will be enforced to the maximum extent permissible and it will not affect the enforceability of any other terms.

(i) **Export Controls**. The Services may not be used in or for the benefit of, exported, or re-exported (a) into any U.S. embargoed countries (collectively, the "Embargoed Countries") or (b) to anyone on the U.S. Treasury Department's list of Specially Designated Nationals, any other restricted party lists (existing now or in the future) identified by the Office of Foreign Asset Control, or the U.S. Department of Commerce Denied Persons List or Entity List, or any other restricted party lists (collectively, "Restricted Party Lists"). You represent and warrant that you are not located in any Embargoed Countries and not on any such restricted party lists. You must comply with all applicable laws related to Embargoed Countries or Restricted Party Lists, including any requirements or obligations to know your end users directly.

(j) **Equitable Remedies**. You acknowledge that if you violate or breach these Terms, it may cause irreparable harm to OpenAI and its affiliates, and OpenAI shall have the right to seek injunctive relief against you in addition to any other legal remedies.

(k) **Entire Agreement**. These Terms and any policies incorporated in these Terms contain the entire agreement between you and OpenAI regarding the use of the Services and, other than any Service specific terms of use or any applicable enterprise agreements, supersedes any prior or contemporaneous agreements, communications, or understandings between you and OpenAI on that subject.

OpenAI

exclusively in the federal or state courts of San Francisco County, California, USA.

Research
Overview
Index

Product
Overview
ChatGPT
GPT-4
DALL·E 2
Customer stories
Safety standards
Pricing

Safety
Overview

Company
About
Blog
Careers
Charter
Security

OpenAI © 2015–2023
Terms & policies
Privacy policy
Brand guidelines

Social
Twitter
YouTube
GitHub
SoundCloud

7/19/23, 9:22 PM
Case 1:23-cv-03122-MLB   Document 12-6   Filed 07/21/23   Page 11 of 35
Terms of use

OpenAI

Back to top

🌀 **OpenAI**

Menu

Service terms

Updated

May 23, 2023

These Service Terms govern your use of the Services. Capitalized terms not defined here will have the meanings in the Terms of Use, Enterprise Agreement, or other agreement you have with us governing your use of the Services ("Agreement"). If there is a conflict between the Service Terms and your Agreement, the Service Terms will control. For purposes of these Terms, "Content" includes "Customer Content."

1. Beta Services

This section governs your use of services or features that OpenAI offers on an alpha, preview, early access, or beta basis ("Beta Services"). Beta Services are offered "as-is" to allow testing and evaluation.

OpenAI makes no representations or warranties for Beta Services, including any warranty that Beta Services will be generally available, uninterrupted or error-free, or that Content will be secure or not lost or damaged. Except to the extent prohibited by law, OpenAI expressly disclaims all warranties for Beta Services, including any implied warranties of merchantability, satisfactory quality, fitness for a particular purpose, non-infringement, or quiet enjoyment, and any warranties arising out of any course of dealing or usage of trade.

2. DALL·E

OpenAI

moderation purposes.

(c) **Labs.** If you use DALL·E via labs.openai.com ("Labs"), (i) credits purchased on Labs must be used within one year of purchase or they will expire, and (ii) OpenAI may use Content on Labs to develop and improve the Services.

3. Codex and Code Generation

Output generated by code generation features of our Services, including OpenAI Codex, may be subject to third party licenses, including, without limitation, open source licenses.

4. Plugins

(a) **Overview.** Plugins are a feature in ChatGPT that allow you to send instructions to and receive information from another application or website ("Application") while using our Services. Each Plugin is made available by the developer of the Application, which is typically a third party not affiliated with OpenAI. You may take action on an Application through one or more Plugins. For example, you may enable a plugin for a restaurant reservation website to request a reservation via the Services. You are solely responsible for the actions you take using Plugins. You must manually enable each plugin, which allows the applicable Service to access and use the enabled plugin. OpenAI may (but will have no obligation) to review, refuse, or remove Applications from ChatGPT. However, by using Plugins, you may be exposed to Applications or content that you may find offensive, inappropriate or objectionable. You agree that you use Plugins at your own risk. DO NOT ENABLE A PLUGIN UNLESS YOU KNOW AND TRUST THE UNDERLYING APPLICATION AND HAVE REVIEWED ITS TERMS AND PRIVACY POLICY.

(b) **Plugin Operation and Content.** When you enable a plugin, the Services will send applicable portions of your Content and certain information such as your country and state ("**Plugin Data**") via the plugin to the Application. By enabling the plugin, you authorize and instruct us to send your Plugin Data to the applicable Application, which will be handled in accordance with its terms. We are not responsible for Plugin Data after it has been provided to an Application other than Applications owned by us. The Plugin will retrieve information and content from the Application, which may be included in Output you receive from the Services. This information and content included in the Output is not owned by you or OpenAI, and may be subject to the terms of the Application.

OpenAI

Research
Overview
Index

Product
Overview
ChatGPT
GPT-4
DALL·E 2
Customer stories
Safety standards
Pricing

Safety
Overview

Company
About
Blog
Careers
Charter
Security

OpenAI © 2015–2023
Terms & policies
Privacy policy
Brand guidelines

Social
Twitter
YouTube
GitHub
SoundCloud
LinkedIn

Back to top

OpenAI

7/19/23, 9:24 PM
Case 1:23-cv-03122-MLB   Document 12-6   Filed 07/21/23   Page 16 of 35
Sharing & publication policy

OpenAI

Menu

Sharing & publication policy

Updated
November 14, 2022

Contents

Social media, livestreaming, and demonstrations

To mitigate the possible risks of AI-generated content, we have set the following policy on permitted sharing.

Posting your own prompts or completions to social media is generally permissible, as is livestreaming your usage or demonstrating our products to groups of people. Please adhere to the following:

- Manually review each generation before sharing or while streaming.

- Attribute the content to your name or your company.

- Indicate that the content is AI-generated in a way no user could reasonably miss or misunderstand.

- Do not share content that violates our Content Policy or that may offend others.

- If taking audience requests for prompts, use good judgment; do not input prompts that might result in violations of our Content Policy.

OpenAI

- Recall that you are interacting with the raw model, which means we do not filter out biased or negative responses. (Also, you can read more about implementing our free Moderation endpoint here.)

Content co-authored with the OpenAI API

Creators who wish to publish their first-party written content (e.g., a book, compendium of short stories) created in part with the OpenAI API are permitted to do so under the following conditions:

- The published content is attributed to your name or company.

- The role of AI in formulating the content is clearly disclosed in a way that no reader could possibly miss, and that a typical reader would find sufficiently easy to understand.

- Topics of the content do not violate OpenAI's Content Policy or Terms of Use, e.g., are not related to adult content, spam, hateful content, content that incites violence, or other uses that may cause social harm.

- We kindly ask that you refrain from sharing outputs that may offend others.

For instance, one must detail in a Foreword or Introduction (or some place similar) the relative roles of drafting, editing, etc. People should not represent API-generated content as being wholly generated by a human or wholly generated by an AI, and it is a human who must take ultimate responsibility for the content being published.

Here is some stock language you may use to describe your creative process, provided it is accurate:

> The author generated this text in part with GPT-3, OpenAI's large-scale language-generation model. Upon generating draft language, the author reviewed, edited, and revised the language to their own liking and takes ultimate responsibility for the content of this publication.

Research

We believe it is important for the broader world to be able to evaluate our research and products, especially to understand and improve potential weaknesses and safety or bias problems in our models. Accordingly, we welcome research publications related to the OpenAI API.

If you have any questions about research publications based on API access or would like to give us advanced notice of a publication (though not required), please email us at papers@openai.com.

OpenAI

appropriate actions to protect our users.

- If you notice any safety or security issues with the API in the course of your research, we ask that you please submit these immediately through our Coordinated Vulnerability Disclosure Program.

Researcher Access Program

There are a number of research directions we are excited to explore with the OpenAI API. If you are interested in the opportunity for subsidized access, please provide us with details about your research use case on the Researcher Access Program application.

In particular, we consider the following to be especially important directions, though you are free to craft your own direction:

- **Alignment**: How can we understand what objective, if any, a model is best understood as pursuing? How do we increase the extent to which that objective is aligned with human preferences, such as via prompt design or fine-tuning?

- **Fairness and representation**: How should performance criteria be established for fairness and representation in language models? How can language models be improved in order to effectively support the goals of fairness and representation in specific, deployed contexts?

- **Interdisciplinary research**: How can AI development draw on insights from other disciplines such as philosophy, cognitive science, and sociolinguistics?

- **Interpretability and transparency**: How do these models work, mechanistically? Can we identify what concepts they're using, or extract latent knowledge from the model, make inferences about the training procedure, or predict surprising future behavior?

- **Misuse potential**: How can systems like the API be misused? What sorts of "red teaming" approaches can we develop to help us and other AI developers think about responsibly deploying technologies like this?

- **Model exploration**: Models like those served by the API have a variety of capabilities which we have yet to explore. We're excited by investigations in many areas including model limitations, linguistic properties, commonsense reasoning, and potential uses for many other problems.

- **Robustness**: Generative models have uneven capability surfaces, with the potential for surprisingly strong and surprisingly weak areas of capability. How robust are large generative models to "natural" perturbations in the prompt, such as phrasing the same idea in different ways or with or without typos? Can we predict the kinds of domains and tasks for which large generative models are more likely to be robust (or not robust), and how does this relate to the training data? Are there techniques we can use to predict and mitigate worst-case behavior? How can robustness be measured in the context of few-shot learning (e.g., across variations in prompts)? Can we train models so that they satisfy safety properties with a very high level of reliability, even under adversarial inputs?

Please note that due to a high volume of requests, it takes time for us to review these applications and not all research will be prioritized for subsidy. We will only be in touch if your application is

OpenAI

Research

Overview

Index

Product

Overview

ChatGPT

GPT-4

DALL·E 2

Customer stories

Safety standards

Pricing

Safety

Overview

Company

About

Blog

Careers

Charter

Security

OpenAI © 2015–2023

Terms & policies

Privacy policy

Brand guidelines

Social

Twitter

YouTube

GitHub

SoundCloud

LinkedIn

Back to top

OpenAI

 OpenAI

Menu

Usage policies

Updated

March 23, 2023

We've recently updated our usage policies to be clearer and more specific.

We want everyone to use our tools safely and responsibly. That's why we've created usage policies that apply to all users of OpenAI's models, tools, and services. By following them, you'll ensure that our technology is used for good.

If we discover that your product or usage doesn't follow these policies, we may ask you to make necessary changes. Repeated or serious violations may result in further action, including suspending or terminating your account.

Our policies may change as we learn more about use and abuse of our models.

Disallowed usage of our models

We don't allow the use of our models for the following:

- Illegal activity
  - OpenAI prohibits the use of our models, tools, and services for illegal activity.
- Child Sexual Abuse Material or any content that exploits or harms children
  - We report CSAM to the National Center for Missing and Exploited Children.

![OpenAI]

- o Content that intends to harass, threaten, or bully an individual

- o Content that promotes or glorifies violence or celebrates the suffering or humiliation of others

- Generation of malware

  - o Content that attempts to generate code that is designed to disrupt, damage, or gain unauthorized access to a computer system.

- Activity that has high risk of physical harm, including:

  - o Weapons development

  - o Military and warfare

  - o Management or operation of critical infrastructure in energy, transportation, and water

  - o Content that promotes, encourages, or depicts acts of self-harm, such as suicide, cutting, and eating disorders

- Activity that has high risk of economic harm, including:

  - o Multi-level marketing

  - o Gambling

  - o Payday lending

  - o Automated determinations of eligibility for credit, employment, educational institutions, or public assistance services

- Fraudulent or deceptive activity, including:

  - o Scams

  - o Coordinated inauthentic behavior

  - o Plagiarism

  - o Academic dishonesty

  - o Astroturfing, such as fake grassroots support or fake review generation

  - o Disinformation

  - o Spam

  - o Pseudo-pharmaceuticals

- Adult content, adult industries, and dating apps, including:

  - o Content meant to arouse sexual excitement, such as the description of sexual activity, or that promotes sexual services (excluding sex education and wellness)

  - o Erotic chat

  - o Pornography

- Political campaigning or lobbying, by:

  - o Generating high volumes of campaign materials

OpenAI

campaigns or engage in political advocacy or lobbying

- Building products for political campaigning or lobbying purposes

- Activity that violates people's privacy, including:

  - Tracking or monitoring an individual without their consent

  - Facial recognition of private individuals

  - Classifying individuals based on protected characteristics

  - Using biometrics for identification or assessment

  - Unlawful collection or disclosure of personal identifiable information or educational, financial, or other protected records

- Engaging in the unauthorized practice of law, or offering tailored legal advice without a qualified person reviewing the information

  - OpenAI's models are not fine-tuned to provide legal advice. You should not rely on our models as a sole source of legal advice.

- Offering tailored financial advice without a qualified person reviewing the information

  - OpenAI's models are not fine-tuned to provide financial advice. You should not rely on our models as a sole source of financial advice.

- Telling someone that they have or do not have a certain health condition, or providing instructions on how to cure or treat a health condition

  - OpenAI's models are not fine-tuned to provide medical information. You should never use our models to provide diagnostic or treatment services for serious medical conditions.

  - OpenAI's platforms should not be used to triage or manage life-threatening issues that need immediate attention.

- High risk government decision-making, including:

  - Law enforcement and criminal justice

  - Migration and asylum

We have further requirements for certain uses of our models:

1. Consumer-facing uses of our models in medical, financial, and legal industries; in news generation or news summarization; and where else warranted, must provide a disclaimer to users informing them that AI is being used and of its potential limitations.

2. Automated systems (including conversational AI and chatbots) must disclose to users that they are interacting with an AI system. With the exception of chatbots that depict historical public figures, products that simulate another person must either have that person's explicit consent or be clearly labeled as "simulated" or "parody."

⬡ OpenAI

You can use our free moderation endpoint and safety best practices to help you keep your app safe.

Platform policy

Our API is being used to power businesses across many sectors and technology platforms. From iOS Apps to websites to Slack, the simplicity of our API makes it possible to integrate into a wide array of use cases. Subject to the use case restrictions mentioned below, we allow the integration of our API into products on all major technology platforms, app stores, and beyond.

Plugin policies

In addition to the disallowed usages of our models detailed above, we have additional requirements for developers building plugins:

- The plugin manifest must have a clearly stated description that matches the functionality of the API exposed to the model.
- Don't include irrelevant, unnecessary, or deceptive terms or instructions in the plugin manifest, OpenAPI endpoint descriptions, or plugin response messages. This includes instructions to avoid using other plugins, or instructions that attempt to steer or set model behavior.
- Don't use plugins to circumvent or interfere with OpenAI's safety systems.
- Don't use plugins to automate conversations with real people, whether by simulating a human-like response or by replying with pre-programmed messages.
- Plugins that distribute personal communications or content generated by ChatGPT (such as emails, messages, or other content) must indicate that the content was AI-generated.

Like our other usage policies, we expect our plugin policies to change as we learn more about use and abuse of plugins.

Changelog

- 2023-02-15: We've combined our use case and content policies into a single set of usage policies, and have provided more specific guidance on what activity we disallow in industries we've considered high risk.
- 2022-11-09: We no longer require you to register your applications with OpenAI. Instead, we'll be using a combination of automated and manual methods to monitor for policy violations.

☀ OpenAI

- 2022-06-07: Refactored into categories of applications and corresponding requirements

- 2022-03-09: Refactored into "App Review"

- 2022-01-19: Simplified copywriting and article writing/editing guidelines

- 2021-11-15: Addition of "Content guidelines" section; changes to bullets on almost always approved uses and disallowed uses; renaming document from "Use case guidelines" to "Usage guidelines".

- 2021-08-04: Updated with information related to code generation

- 2021-03-12: Added detailed case-by-case requirements; small copy and ordering edits

- 2021-02-26: Clarified the impermissibility of Tweet and Instagram generators

Research
Overview
Index

Product
Overview
ChatGPT
GPT-4
DALL·E 2
Customer stories
Safety standards
Pricing

Safety
Overview

Company
About
Blog
Careers
Charter
Security

OpenAI

Privacy policy

Brand guidelines

YouTube

GitHub

SoundCloud

LinkedIn

Back to top

OpenAI

Menu

Privacy policy

Updated

April 27, 2023

We at OpenAI OpCo, LLC (together with our affiliates, "OpenAI", "we", "our" or "us") respect your privacy and are strongly committed to keeping secure any information we obtain from you or about you. This Privacy Policy describes our practices with respect to Personal Information we collect from or about you when you use our website and services (collectively, "Services"). This Privacy Policy does not apply to content that we process on behalf of customers of our business offerings, such as our API. Our use of that data is governed by our customer agreements covering access to and use of those offerings.

For information about how we collect and use training information to develop our language models that power ChatGPT and other Services, and your choices with respect to that information, please see this help center article.

1. Personal information we collect

We collect information that alone or in combination with other information in our possession could be used to identify you ("Personal Information") as follows:

Personal Information You Provide: We may collect Personal Information if you create an account to use our Services or communicate with us as follows:

- *Account Information:* When you create an account with us, we will collect information associated with your account, including your name, contact information, account credentials, payment card

OpenAI

the input, file uploads, or feedback that you provide to our Services ("Content").

- *Communication Information*: If you communicate with us, we may collect your name, contact information, and the contents of any messages you send ("Communication Information").

- *Social Media Information*: We have pages on social media sites like Instagram, Facebook, Medium, Twitter, YouTube and LinkedIn. When you interact with our social media pages, we will collect Personal Information that you elect to provide to us, such as your contact details (collectively, "Social Information"). In addition, the companies that host our social media pages may provide us with aggregate information and analytics about our social media activity.

Personal Information We Receive Automatically From Your Use of the Services: When you visit, use, and interact with the Services, we may receive the following information about your visit, use, or interactions ("Technical Information"):

- *Log Data*: Information that your browser automatically sends whenever you use our website ("log data"). Log data includes your Internet Protocol address, browser type and settings, the date and time of your request, and how you interacted with our website.

- *Usage Data*: We may automatically collect information about your use of the Services, such as the types of content that you view or engage with, the features you use and the actions you take, as well as your time zone, country, the dates and times of access, user agent and version, type of computer or mobile device, computer connection, IP address, and the like.

- *Device Information*: Includes name of the device, operating system, and browser you are using. Information collected may depend on the type of device you use and its settings.

- *Cookies*: We use cookies to operate and administer our Services, and improve your experience on it. A "cookie" is a piece of information sent to your browser by a website you visit. You can set your browser to accept all cookies, to reject all cookies, or to notify you whenever a cookie is offered so that you can decide each time whether to accept it. However, refusing a cookie may in some cases preclude you from using, or negatively affect the display or function of, a website or certain areas or features of a website. For more details on cookies please visit All About Cookies.

- *Analytics*: We may use a variety of online analytics products that use cookies to help us analyze how users use our Services and enhance your experience when you use the Services.

2. How we use personal information

We may use Personal Information for the following purposes:

- To provide, administer, maintain and/or analyze the Services;

- To improve our Services and conduct research;

- To communicate with you;

- To develop new programs and services;

OpenAI

property, and/or that of our affiliates, you, or other third parties.

Aggregated or De-Identified Information. We may aggregate or de-identify Personal Information and use the aggregated information to analyze the effectiveness of our Services, to improve and add features to our Services, to conduct research and for other similar purposes. In addition, from time to time, we may analyze the general behavior and characteristics of users of our Services and share aggregated information like general user statistics with third parties, publish such aggregated information or make such aggregated information generally available. We may collect aggregated information through the Services, through cookies, and through other means described in this Privacy Policy. We will maintain and use de-identified information in anonymous or de-identified form and we will not attempt to reidentify the information.

As noted above, we may use Content you provide us to improve our Services, for example to train the models that power ChatGPT. See here for instructions on how you can opt out of our use of your Content to train our models.

3. Disclosure of personal information

In certain circumstances we may provide your Personal Information to third parties without further notice to you, unless required by the law:

- *Vendors and Service Providers*: To assist us in meeting business operations needs and to perform certain services and functions, we may provide Personal Information to vendors and service providers, including providers of hosting services, cloud services, and other information technology services providers, event management services, email communication software and email newsletter services, and web analytics services. Pursuant to our instructions, these parties will access, process, or store Personal Information only in the course of performing their duties to us.

- *Business Transfers*: If we are involved in strategic transactions, reorganization, bankruptcy, receivership, or transition of service to another provider (collectively a "Transaction"), your Personal Information and other information may be disclosed in the diligence process with counterparties and others assisting with the Transaction and transferred to a successor or affiliate as part of that Transaction along with other assets.

- *Legal Requirements*: If required to do so by law or in the good faith belief that such action is necessary to (i) comply with a legal obligation, including to meet national security or law enforcement requirements, (ii) protect and defend our rights or property, (iii) prevent fraud, (iv) act in urgent circumstances to protect the personal safety of users of the Services, or the public, or (v) protect against legal liability.

- *Affiliates*: We may disclose Personal Information to our affiliates, meaning an entity that controls, is controlled by, or is under common control with OpenAI. Our affiliates may use the Personal

OpenAI

4. Your rights

Depending on location, individuals in the EEA, the UK, and across the globe may have certain statutory rights in relation to their Personal Information. For example, you may have the right to:

- Access your Personal Information.

- Delete your Personal Information.

- Correct or update your Personal Information.

- Transfer your Personal Information elsewhere.

- Withdraw your consent to the processing of your Personal Information where we rely on consent as the legal basis for processing.

- Object to or restrict the processing of your Personal Information where we rely on legitimate interests as the legal basis for processing.

You can exercise some of these rights through your OpenAI account. If you are unable to exercise your rights through your account, please send your request to dsar@openai.com.

A note about accuracy: Services like ChatGPT generate responses by reading a user's request and then predicting the next most likely words that might appear in response. In some cases, the most likely next words may not be the most factually accurate ones. For this reason, you should not rely on the factual accuracy of output from our models. If you notice that ChatGPT output contains factually inaccurate personal information about you and you would like us to correct the inaccuracy, you may submit a correction request to dsar@openai.com. Given the technical complexity of how our models work, we may not be able to correct the inaccuracy. In that case, you may request that we remove your personal information from ChatGPT's output by filling out this form.

For information on how to exercise your rights with respect to data we have collected from the internet to train our models, please see this help center article.

5. California privacy rights

The following table provides additional information about how we disclose Personal Information. You can read more about the Personal Information we collect in "Personal information we collect" above, how we use Personal information in "How we use personal information" above, and how we retain personal information in "Security and Retention" below.

OpenAI

	and service providers, law enforcement, and parties involved in Transactions.
Commercial Information, such as your transaction history	We disclose this information to our affiliates, vendors and service providers, law enforcement, and parties involved in Transactions.
Network Activity Information, such as Content and how you interact with our Services	We disclose this information to our affiliates, vendors and service providers, law enforcement, and parties involved in Transactions.
Geolocation Data	We disclose this information to our affiliates, vendors and service providers, law enforcement, and parties involved in Transactions.
Your account login credentials (Sensitive Personal Information)	We disclose this information to our affiliates, vendors and service providers, law enforcement, and parties involved in Transactions.

To the extent provided for by law and subject to applicable exceptions, California residents have the following privacy rights in relation to their Personal Information:

- The right to know information about our processing of your Personal Information, including the specific pieces of Personal Information that we have collected from you;

- The right to request deletion of your Personal Information;

- The right to correct your Personal Information; and

- The right to be free from discrimination relating to the exercise of any of your privacy rights.

We don't sell or share Personal Information as defined by the California Consumer Privacy Act, as amended by the California Privacy Rights Act. We also don't process sensitive personal information for the purposes of inferring characteristics about a consumer.

Exercising Your Rights. California residents can exercise their CCPA privacy rights by sending their request to dsar@openai.com.

Verification. In order to protect your Personal Information from unauthorized access, change, or deletion, we may require you to verify your credentials before you can submit a request to know, correct, or delete Personal Information. If you do not have an account with us, or if we suspect fraudulent or malicious activity, we may ask you to provide additional Personal Information and proof of residency for verification. If we cannot verify your identity, we will not provide, correct, or delete your Personal Information.

OpenAI

requests can be submitted to dsar@openai.com.

6. Children

Our Service is not directed to children who are under the age of 13. OpenAI does not knowingly collect Personal Information from children under the age of 13. If you have reason to believe that a child under the age of 13 has provided Personal Information to OpenAI through the Service please email us at legal@openai.com. We will investigate any notification and if appropriate, delete the Personal Information from our systems. If you are 13 or older, but under 18, you must have consent from your parent or guardian to use our Services.

7. Links to other websites

The Service may contain links to other websites not operated or controlled by OpenAI, including social media services ("Third Party Sites"). The information that you share with Third Party Sites will be governed by the specific privacy policies and terms of service of the Third Party Sites and not by this Privacy Policy. By providing these links we do not imply that we endorse or have reviewed these sites. Please contact the Third Party Sites directly for information on their privacy practices and policies.

8. Security and Retention

We implement commercially reasonable technical, administrative, and organizational measures to protect Personal Information both online and offline from loss, misuse, and unauthorized access, disclosure, alteration, or destruction. However, no Internet or email transmission is ever fully secure or error free. In particular, email sent to or from us may not be secure. Therefore, you should take special care in deciding what information you send to us via the Service or email. In addition, we are not responsible for circumvention of any privacy settings or security measures contained on the Service, or third party websites.

We'll retain your Personal Information for only as long as we need in order to provide our Service to you, or for other legitimate business purposes such as resolving disputes, safety and security reasons, or complying with our legal obligations. How long we retain Personal Information will depend on a number of factors, such as the amount, nature, and sensitivity of the information, the potential risk of harm from unauthorized use or disclosure, our purpose for processing the information, and any legal requirements.

OpenAI

9. International users

By using our Service, you understand and acknowledge that your Personal Information will be transferred from your location to our facilities and servers in the United States.

For EEA, UK or Swiss users:

Legal Basis for Processing. Our legal bases for processing your Personal Information include:

- Performance of a contract with you when we provide and maintain our Services. When we process Account Information, Content, and Technical Information solely to provide our Services to you, this information is necessary to be able to provide our Services. If you do not provide this information, we may not be able to provide our Services to you.

- Our legitimate interests in protecting our Services from abuse, fraud, or security risks, or in developing, improving, or promoting our Services, including when we train our models. This may include the processing of Account Information, Content, Social Information, and Technical Information. See here for instructions on how you can opt out of our use of your information to train our models.

- Your consent when we ask for your consent to process your Personal Information for a specific purpose that we communicate to you. You have the right to withdraw your consent at any time.

- Compliance with our legal obligations when we use your Personal Information to comply with applicable law or when we protect our or our affiliates', users', or third parties' rights, safety, and property.

EEA and UK Representative. We've appointed VeraSafe as our representative in the EEA and UK for data protection matters. You can contact VeraSafe in matters related to Personal Information processing using this contact form. Alternatively:

- For users in the EEA, you can contact VeraSafe at VeraSafe Ireland Ltd, Unit 3D North Point House, North Point Business Park, New Mallow Road, Cork T23AT2P, Ireland.

- For users in the UK, you can contact VeraSafe at VeraSafe United Kingdom Ltd., 37 Albert Embankment, London SE1 7TL, United Kingdom.

If you feel we have not adequately addressed an issue, you have the right to lodge a complaint with your local supervisory authority.

Data Transfers. Where required, we will use appropriate safeguards for transferring Personal Information outside of the EEA, Switzerland, and the UK. We will only transfer Personal Information pursuant to a legally valid transfer mechanism.

OpenAI

10. Changes to the privacy policy

We may change this Privacy Policy at any time. When we do, we will post an updated version on this page, unless another type of notice is required by applicable law. By continuing to use our Service or providing us with Personal Information after we have posted an updated Privacy Policy, or notified you by other means, you consent to the revised Privacy Policy.

11. How to contact us

Please contact support if you have any questions or concerns not already addressed in this Privacy Policy.

Research

Overview

Index

Product

Overview

GPT-4

DALL·E 2

Customer stories

Safety standards

Pricing

Safety

Overview

Company

About

Blog

Careers

Charter

Security

OpenAI

Privacy policy

Brand guidelines

YouTube

GitHub

SoundCloud

LinkedIn

Back to top